

CLAIMS

We claim:

1. A method of managing communications between service components in a cluster-based computing environment, the cluster-based computing environment comprising a plurality of processing nodes interconnected via a network switching system, each of the service components being programmed on a respective one of the processing nodes, the method comprising:

configuring filter logic in the cluster-based computing environment with rules representative of allowed inter-node communications between service components;

detecting an attempted inter-node communication between service components;

applying the filter logic to determine that the attempted inter-node communication is not allowed; and

responsively blocking the attempted inter-node communication.

2. The method of claim 1,

wherein each of the service components is associated with a respective service-access-point (SAP) in the cluster-based computing environment; and

wherein configuring the filter logic with rules representative of allowed inter-node communications between service components comprises configuring the filter logic with rules indicating allowed communications between respective SAPs in the cluster-based computing environment.

3. The method of claim 2, wherein the respective SAP of each service component comprises an Internet Protocol address of the respective processing node on which the service component is programmed.

4. The method of claim 3, wherein at least one of the SAPs further comprises a port selected from the group consisting of a TCP port and a UDP port.

5. The method of claim 1,
wherein the communications between service components are packet-based; and
wherein configuring the filter logic with rules representative of allowed inter-node communications between service components comprises configuring the filter logic with rules each indicating an allowed combination of at least (i) a packet transport protocol, (ii) a source address in the cluster-based computing environment and (iii) a destination address in the cluster-based computing environment.

6. The method of claim 1, wherein the network switching system comprises a switch, and wherein configuring filter logic in the cluster-based computing environment comprises setting up the switch to enforce the rules.

7. The method of claim 6,
wherein the switch comprises a packet-filtering agent and a provisioning-interface through which instructions may be provided to set up the packet-filtering agent, the switch being

arranged to translate the instructions into packet-filtering logic executable by the packet-filtering
5 agent; and

wherein setting up the switch to enforce the rules comprises providing the switch, via the
provisioning-interface, with instructions representative of the rules.

8. The method of claim 6, wherein:

each inter-node communication comprises a packet including a VLAN tag; and

setting up the switch to enforce the rules comprises setting up the switch with VLAN
logic associating each service component with a respective VLAN tag, whereby the switch may
allow a given packet to be routed to a given service component only if the VLAN tag of the
given packet is associated with the given service component.

9. The method of claim 6, wherein setting up the switch to enforce the rules
comprises setting up the switch with a plurality of static packet routes, each static packet route
correlating a respective source service-access-point with a respective destination service-access-
point.

10. The method of claim 1,

wherein at least a given one of the processing nodes includes a firewall for restricting
communications with the given processing node; and

wherein configuring filter logic in the cluster-based computing environment comprises
5 provisioning the firewall of the given processing node to allow communications between at least

one service component programmed on the given processing node and at least one service component programmed on another processing node.

11. The method of claim 1, wherein the attempted inter-node communication comprises an attempted inter-node communication between antagonistic service components.

12. The method of claim 1, wherein:
the attempted inter-node communication comprises an attempted communication of a packet comprised of data; and

applying the filter logic to determine that the attempted inter-node communication is not allowed comprises using the data of the packet to determine that the inter-node communication is not allowed.

13. The method of claim 12, wherein the data represents information selected from the group consisting of (i) source, (ii) destination and (iii) service level.

14. The method of claim 12, wherein using the data of the packet to determine that the attempted inter-node communication is not allowed comprises determining, based at least in part on the data of the packet, that the attempted inter-node communication does not satisfy any of the rules representative of allowed inter-node communications between service components.

15. The method of claim 1, wherein the attempted inter-node communication comprises an attempted communication of a packet from a first processing node to a second

processing node, and wherein blocking the attempted communication comprises dropping the packet.

16. A method for managing application logic in a public computing platform, the public computing platform comprising a network of processing nodes interconnected by a switching system, the method comprising:

receiving specifications of at least two computer-program applications, the applications cooperatively comprising a number of application components;

generating access control rules defining allowed communications between the application components;

loading the application components of the at least two applications onto at least two of the processing nodes of the public computing platform, whereby the processing nodes may then execute the application components; and

provisioning the public computing platform to allow inter-node communications comprising the allowed communications between application components and to disallow other inter-node communications,

whereby, in response to an attempted communication between application components, the public computing platform may determine that the attempted communication is not allowed and may responsively block the attempted communication.

17. The method of claim 16, further comprising:

assigning to each application component a respective trustworthiness measure and a respective criticality measure, and using the trustworthiness and criticality measures of a given application component to select a given processing node of the public computing platform onto
5 which the given application component should be loaded,

wherein loading the application components onto the processing nodes of the public computing platform comprises loading the given application component onto the given processing node.

18. A public computing platform comprising:

a network switching system;

a plurality of processing nodes interconnected via the network switching system;

a plurality of application components loaded onto the processing nodes, each application component having a respective service-access-point defining (i) a network address of the processing node on which the application component is loaded and (ii) a port at the processing node, the port being associated with the application component;

logic indicating allowed inter-node communications between application components;

the logic being executable, in response to an attempted inter-node communication, to

10 make a determination of whether the attempted inter-node communication is allowed; and

the logic being executable, in response to a determination that the attempted inter-node communication is not allowed, to block the attempted inter-node communication.

19. The public computing platform of claim 18, wherein the switching system comprises a network switch having a packet-filtering agent, and wherein the logic is embodied at least in part in the packet-filtering agent.

20. The public computing platform of claim 18, wherein, for a given attempted inter-node communication between a first application component and a second application component, the logic is executable to determine whether the given attempted inter-node communication is allowed based at least in part on a parameter selected from the group consisting of (i) the service-access-point of the first service application component and (ii) the service-access-point of the second application component.

21. The public computing platform of claim 18,
wherein at least a first one of the application components loaded onto a first processing node of the public computing platform is owned by a first application provider, and at least a second one of the application components loaded onto a second processing node of the public computing platform is owned by a second application provider; and

wherein the first application provider is in competition with the second application provider for business.

22. The public computing platform of claim 21, wherein the first application component is a component of a first application, the second application component is a

component of a second application, and the first application and second application are competing applications.

09723558-120400